3.4 A Scalable X86 CPU Design for 90nm Process

J. Schutz, C. Webb

Intel, Hillsboro, OR

This paper describes the design of a third generation [1][2] Pentium 4 processor to meet the challenges of a 90nm technology [3]. The design has 1MB L2 cache, 16kB L1 cache, 125M transistors and 112mm^2 die size. A die photo is shown in Fig. 3.4.1. To increase the efficiency and performance of hyper-threading technology, two new instructions are added. Other performance improvements, such as better lock protocol and larger caches allow the increase of hyper-threading performance compared to previous implementations. Hyper-threading also helps the processor performance to better scale with frequency, as it provides a mechanism for useful work to be performed when the processor waits longer for memory latency.

Designing for a 90nm process presents many challenges. The leakage currents of the transistors have been increasing with each process generation [4]. Design must allow for changes between process generations and be scalable with more improvement in transistors over the life of the process. Interconnect delays do not scale as well as transistors [5]. Test and debug of a high frequency processor is difficult [6] and requires improved design-for-test. On the other hand, new features increase transistor count and design complexity. Better design automation techniques are needed to convert the layout to the new process.

Transistor delays have improved by 30% from the previous process, but interconnect delays have less improvement. Initially, the frequency of a design on a new process may not be RC limited but as transistors improve, the frequency becomes RC-limited. Designing for a new process and for scaling through the process life without becoming RC-limited requires:

- Re-pipeling of multiple clock paths with large RC delays
- Repeater spacing for process end of life (EOL)
- Timing analysis to identify beginning and EOL paths
- Designing arrays to be functional at process EOL
- High frequency clock distribution [7]
- New circuit techniques in integer unit
- Changes to micro-architecture to use more simple circuits

Timing analysis can be done with the beginning and end of life process parameters to find the EOL RC-limited paths, but this requires multiple timing runs. A simpler option changes the path ordering, so that the early transistor-dominated paths and the later RC-dominated paths are apparent in one timing run. Transistor parameters for the early process and 30% increased interconnect resistance are used for timing analysis. This changes the timing path ordering in a single run to make the RC-limited paths appear to be worse as shown in Fig. 3.4.2. Transistor-limited paths are analyzed correctly for the early process. RC-limited paths that will be a problem later in the process life, show up as problems also in the single timing analysis. Word line and bit line signals in register files, PLA, and ROMs need to be designed to have adequate margin for reads and writes for the full range of the process. A combination of the R scaling and simulations with a fast process skew is used to insure that the signals reach full levels at process EOL. The bit line segmentation for most of the register files is changed from 16 bits on local bit line to 8 bits. Figure 3.4.3 shows word line, local and global bit line signals at 6GHz.

The design uses low voltage swing (LVS) technology in t he double-frequency integer core [1]. The circuit technology is described in [8].

The leakage current of the transistor increases to achieve the best performance on 90nm process. The impact of the increased leakage current is minimized by using devices with longer channel length and also careful use of dual-V_T transistors. The bits in the small signal arrays use transistors that are 20 to 30nm longer than minimum. Transistors with 10nm longer than minimum channel length are used in the register files and in non-critical circuits to reduce the leakage power. The process provides transistors with two V_Ts. The lower V_T usage is controlled to insure that these devices are only in the speed-limiting paths and only add small additional leakage current. Low V_T transistors are less than 1% of the total non-memory transistor width. The active and worst case leakage power percentages are shown in Fig. 3.4.4. The power distribution is non-uniform as shown in Fig. 3.4.5. The temperature profile generated from the power distribution is used in the leakage calculations and design for electro-migration and self heating. Additional VSS and VCC routing is done in high-power areas of the die.

The design goal was to use standard cells and to automate the design and layout as much as possible. Figure 3.4.6 shows areas of the die that use cell-based design (CBD) and areas with custom layout. The die has 20 CBD megablocks of 25000 to 100,000 cells. Megablocks have flat layout hierarchy of custom embedded blocks and standard cells. Placement of cells is automated based upon timing and routing requirements. Cells are allowed to spread out around custom blocks as needed for timing. Less than 20% of the cells come from schematics. All other CBD logic is synthesized. Transistor sizing is automated and includes consideration of timing and power. Special tools are developed for automation of clock and mininum-delay buffer insertion and scan routing.

The design includes many design-for-test features to improve debug and manufacturing test. These include extensive use of scan for debug and test, signature mode, built-in pattern generation for test of large caches (PBIST), direct access testing (DAT) test mode, on-die clock shrinking, locate critical path (LCP), within die process monitors, IO loop back and IO test generator (IBIST). The clock distribution is divided into several LCP domains as shown in Fig. 3.4.7. The LCP feature allows clock rising and falling edges in 64 domains to be adjusted with four settings with a total delay of approximately 1.5 inverters. LCP settings can be adjusted by loading settings through the TAP port or with fuses. This allows easy adjustments for more timing margin for sequential set or valid timings. LCP and on-die clock shrinking are used to isolate many critical timings during debug.

References:

[1] D. Sager et al., "A 0.18μm CMOS IA32 Microprocessor with a 4GHz Integer Execution Unit," *ISSCC Dig. Tech. Papers*, pp. 324-325, Feb. 2001.
[2] D. Deleganes et al., "Designing a 3GHz, 130nm, Intel Pentium 4 Processor," *Symp. VLSI Circuits Dig.*, pp. 130-133, Jun. 2002.
[3] S. Thompson, "A 90nm Technology Featuring 50nm Strained Silicon Channel Transistors, 7 Layers of Interconnect, Low K ILD and 1μm^2 Memory Bit," *IEDM Tech. Dig*., pp. 61-64, Dec. 2002.
[4] S. Borkar, "Low-Power Design Challenges for the Decade," *Proc. of ASP-DAC,* pp. 293-296, Feb. 2001.
[5] J. Davis et al., "Interconnect Limits in Gigascale Integration in the 21st Century," *Proceedings of the IEEE*, March 2001
[6] D. D. Josephson, "The Manic Depression of Microprocessor Debug," *Proc. of ITC Conf.,* pp. 657-663, Oct. 2002.
[7] N. Bindal et al., "Scalable Sub-10ps Skew Global Clock Distribution for a 90nm Multi-Ghz IA Microprocessor," *ISSCC Dig. Tech. Papers*, pp. 346-347, Feb. 2003.
[8] D. Deleganes et al., "Low-Voltage-Swing Logic Circuits for a 7GHz x86 Integer Core," *ISSCC Dig. Tech. Papers,* pp. 154, Feb. 2004.

 0-7803-8267-6/04

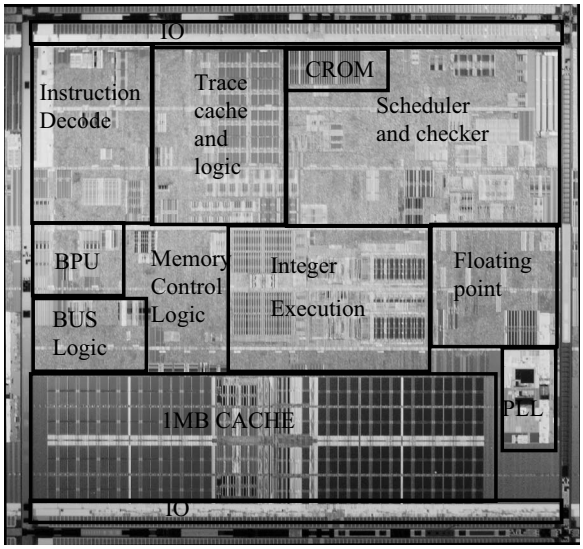

Figure 3.4.1: Chip micrograph.

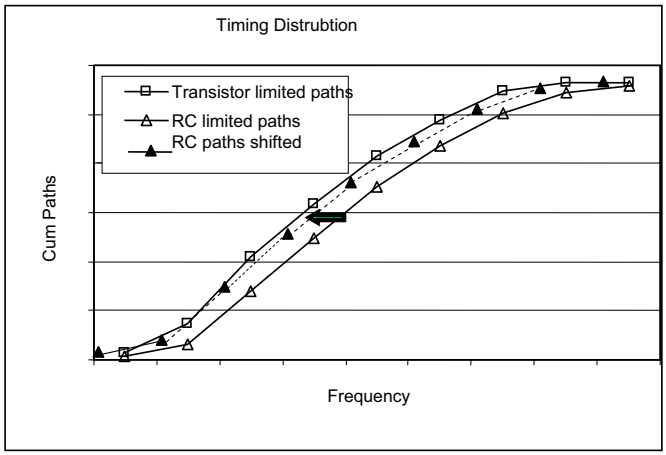

Figure 3.4.2: Timing distribution.

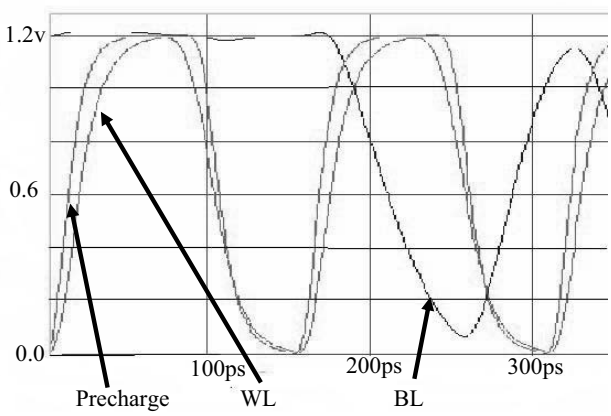

Figure 3.4.3: Word line, bit line and global line signals at 6GHz.

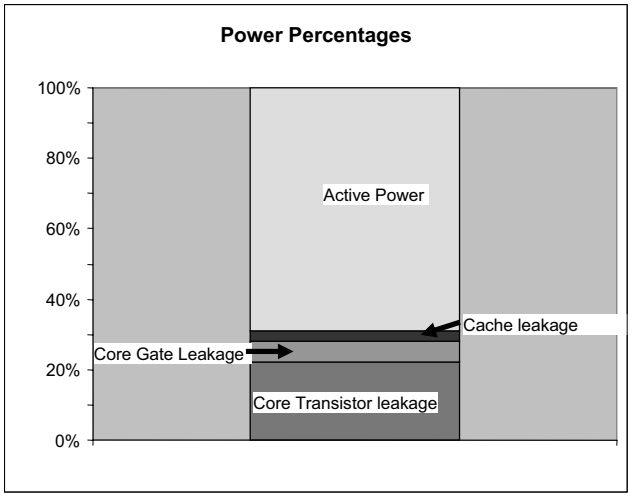

Figure 3.4.4: Active and worst-case leakage power percentages.

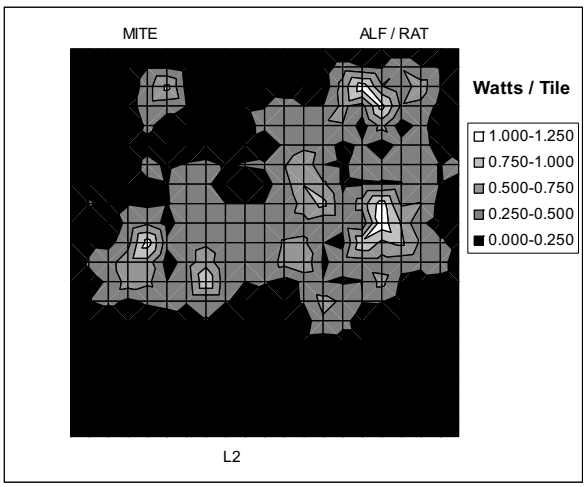

Figure 3.4.5: Power distribution.

Figure 3.4.6: CBD (white) and areas with custom layout (black).

0-7803-8267-6/04

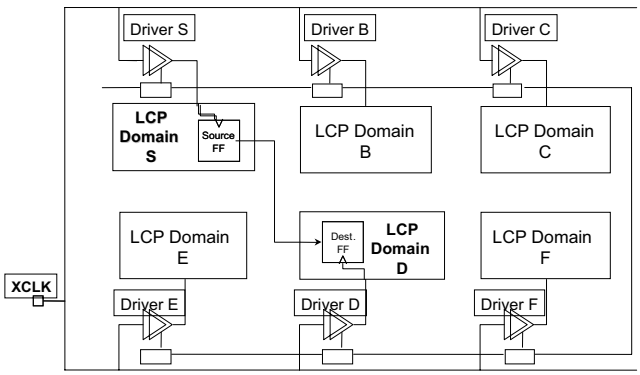

Figure 3.4.7: Clock distribution.

0-7803-8267-6/04

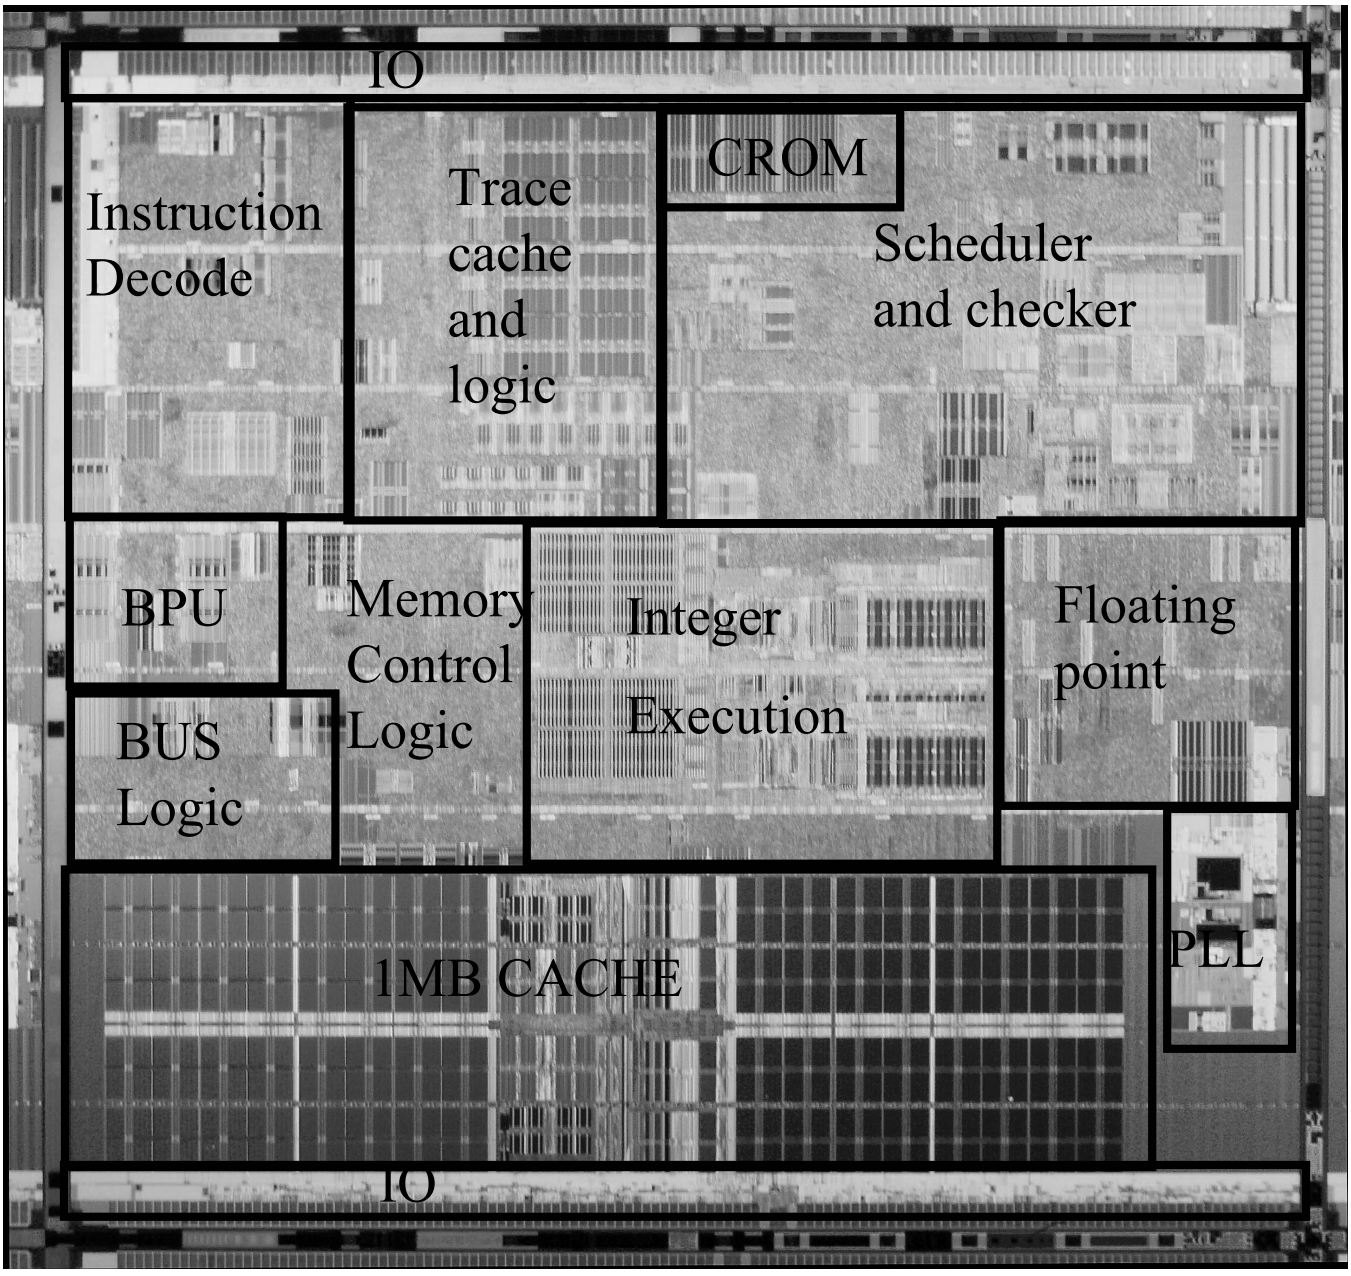

Figure 3.4.1: Chip micrograph.

0-7803-8267-6/04

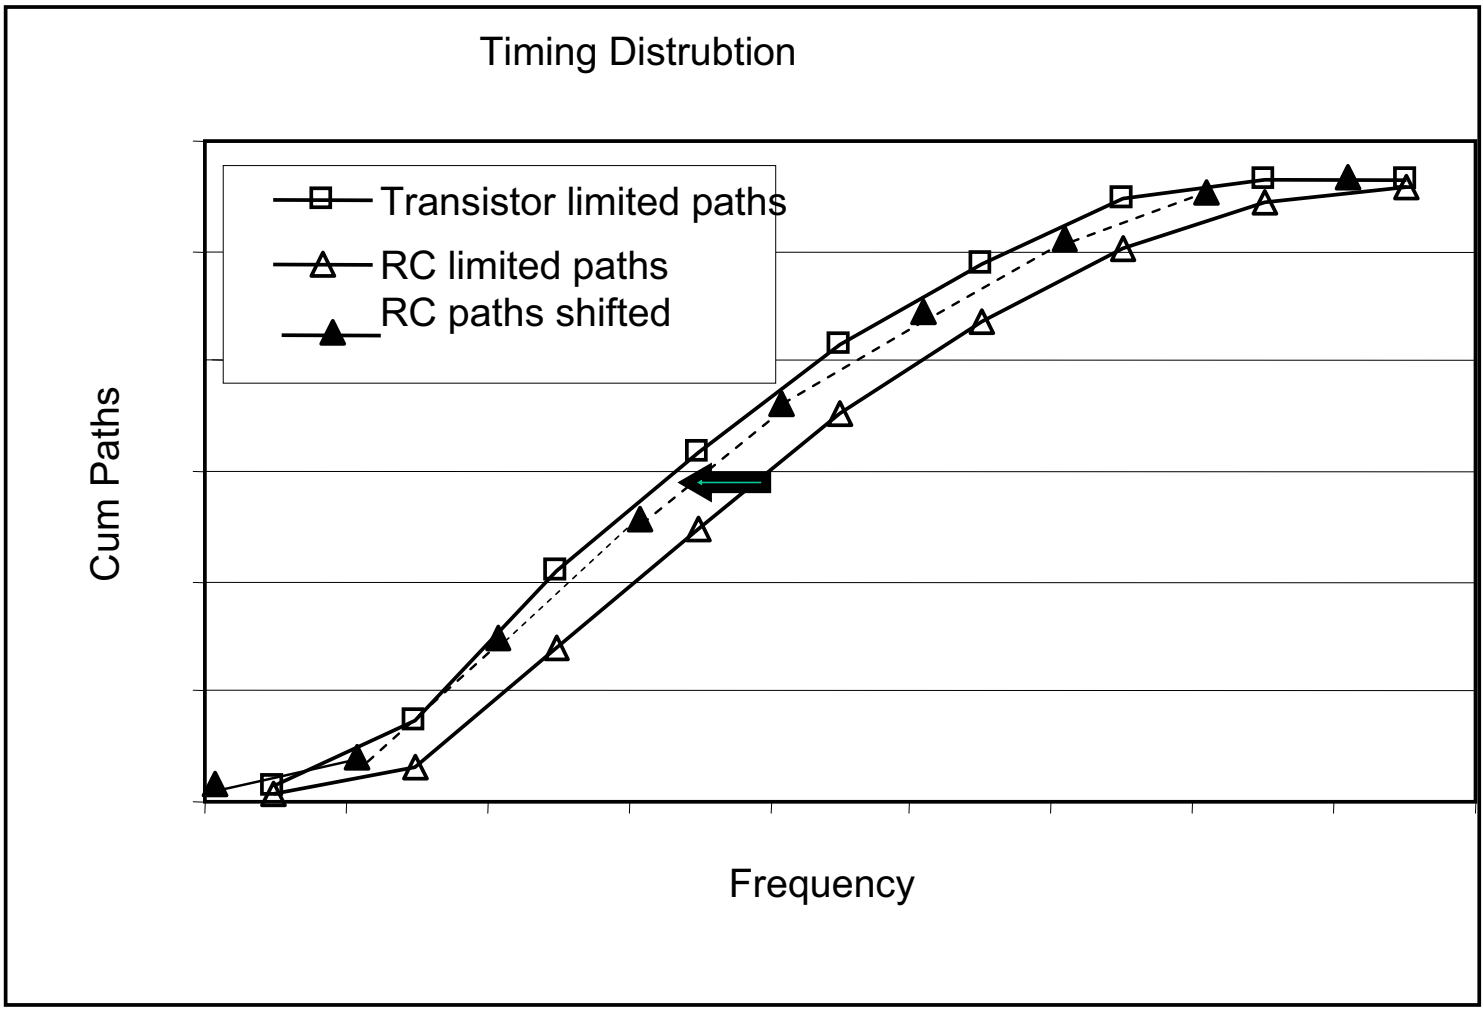

Figure 3.4.2: Timing distribution.

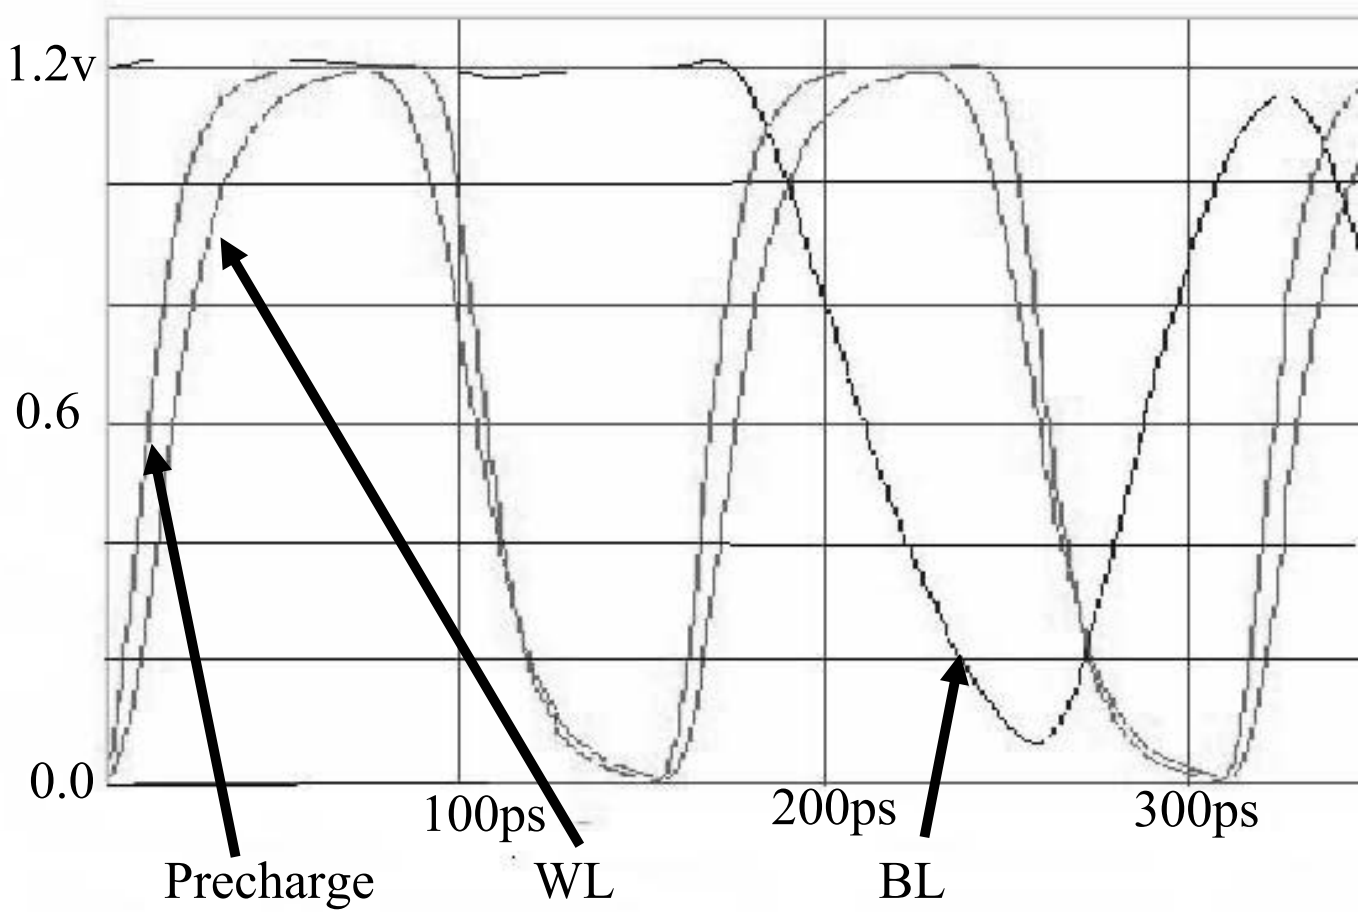

Figure 3.4.3: Word line, bit line and global line signals at 6GHz.

0-7803-8267-6/04

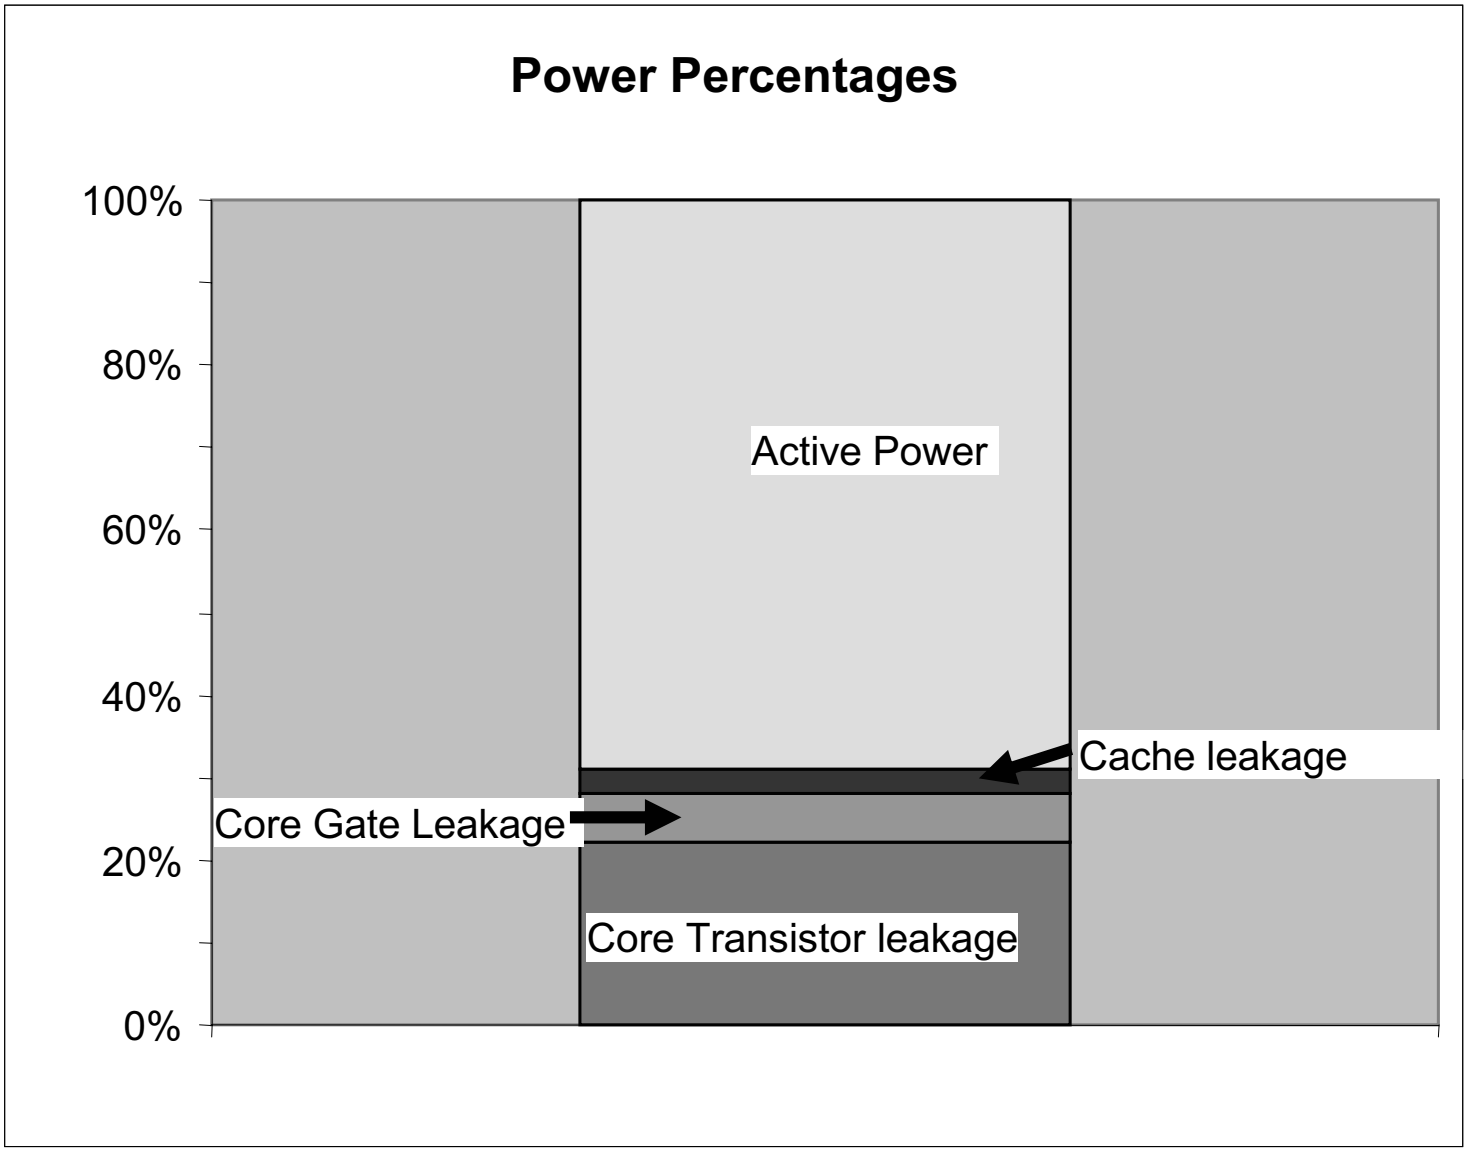

Figure 3.4.4: Active and worst-case leakage power percentages.

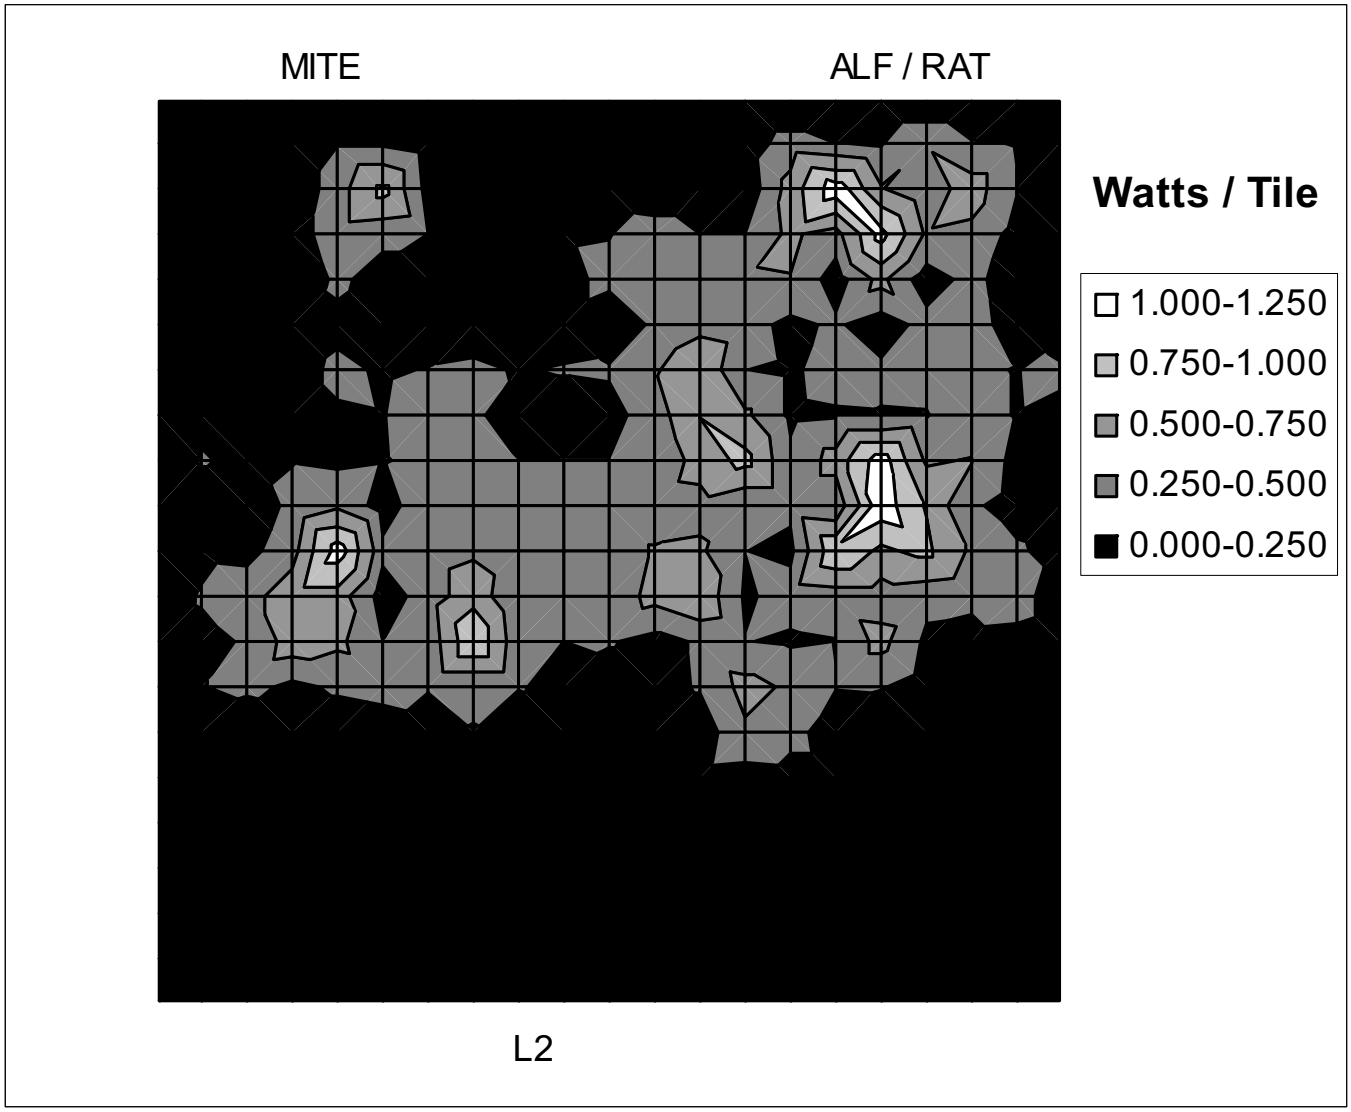

Figure 3.4.5: Power distribution.

Figure 3.4.6: CBD (white) and areas with custom layout (black).

0-7803-8267-6/04

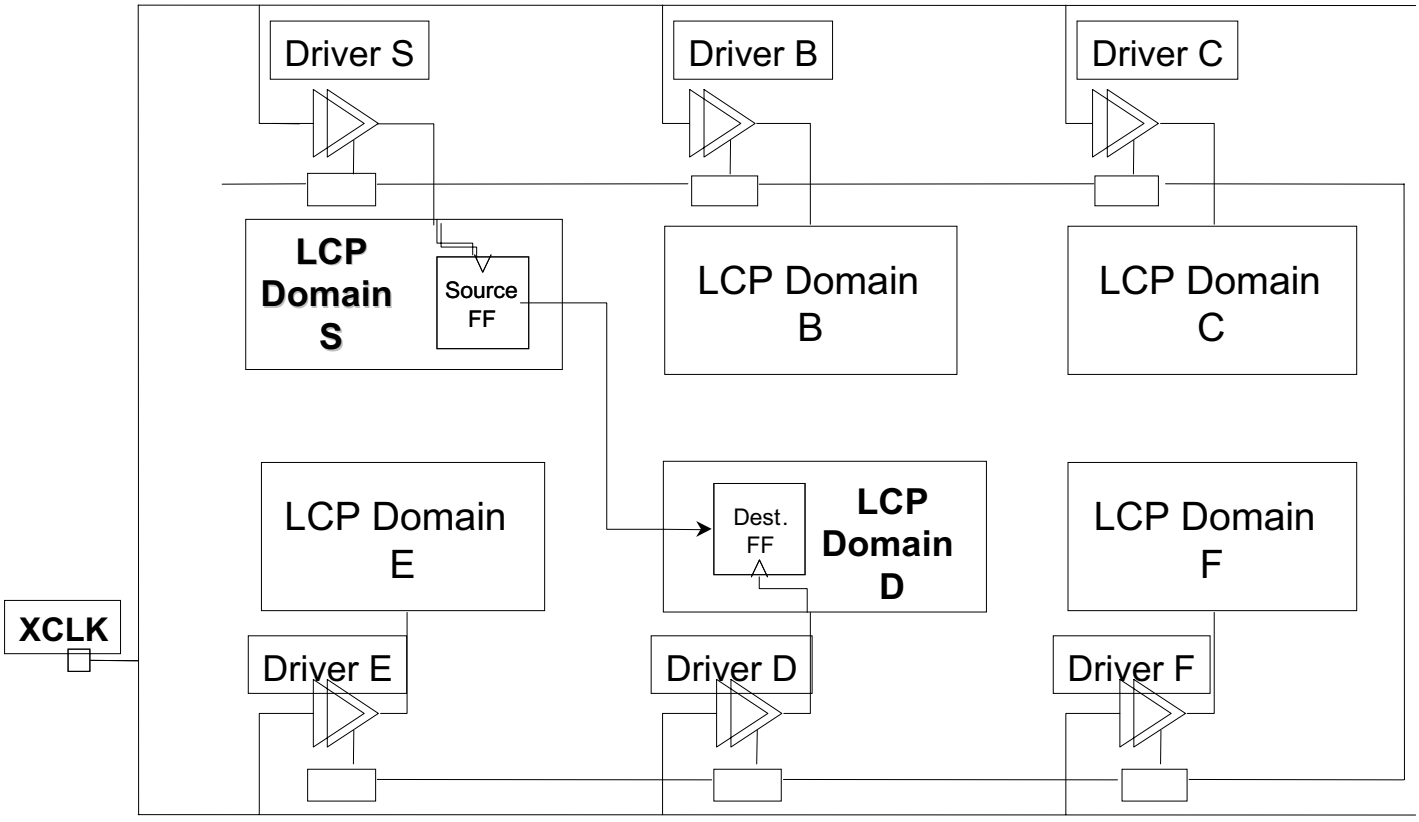

Figure 3.4.7: Clock distribution.